

Bayesian Mode Regression

Keming Yu*

Brunel University, UK

Katerina Aristodemou

Brunel University, UK

Summary. Like mean, quantile and variance, mode is also an important measure of central tendency and data summary. Many practical questions often focus on “Which element (gene or file or signal) occurs most often or is the most typical among all elements in a network?”. In such cases mode regression provides a convenient summary of how the regressors affect the conditional mode and is totally different from other regression models based on conditional mean or conditional quantile or conditional variance. Some inference methods have been used for mode regression but none of them from the Bayesian perspective. This paper introduces Bayesian mode regression by exploring three different approaches. We start from a parametric Bayesian model by employing a likelihood function that is based on a mode uniform distribution. It is shown that irrespective of the original distribution of the data, the use of this special uniform distribution is a very natural and effective way for Bayesian mode regression. Posterior estimates based on this parametric likelihood, even under misspecification, are consistent and asymptotically normal. We then develop a nonparametric Bayesian model by using Dirichlet process (DP) mixtures of mode uniform distributions and finally we explore Bayesian empirical likelihood mode regression by taking empirical likelihood into a Bayesian framework. The paper also demonstrates that a variety of improper priors for the unknown model parameters yield a proper joint posterior. The proposed approach is illustrated using simulated datasets and a real data set.

Keywords: Bayesian inference; empirical likelihood; Markov Chain Monte Carlo methods; mode regression; nonparametric Bayesian; parametric Bayesian.

*Address for correspondence: Department of Mathematical Sciences, Brunel University, UK.

E-mail: keming.yu@brunel.ac.uk

1 Introduction

Mode, the most likely value of a distribution, has wide applications in biology, astronomy, economics and finance. For example, it is not uncommon in many fields to encounter data distributions that are skewed or contain outliers. In those cases, the arithmetic mean may not be an appropriate statistic to represent the center of location of the data. Alternative statistics with less bias are the median and the mode. The mean or median of two densities may be identical, while the shapes of the two densities are quite different. Mode preserves some of the important features, such as wiggles, of the underlying distribution function, whereas the mean or median tend to average out the data. Actually, as an important statistic, mode has been used in modern science to identify the most frequent or the most typical element in certain network systems (Hedges and Shah (2003), Heckman, Geiser, Eidell, Stauffer, Kardos, and Hedges (2001), Kumar and Hedges (1998), Markov et al. (1997)).

A mode estimator is often defined as the maximum of the estimated distribution density, typically under nonparametric kernel estimation. Such mode estimation has attracted a lot of attention in the statistics literature for decades by various authors (Yasukawa (1926), Parzen (1962), Grenander (1965), Eddy (1980), Bickel and Fan (1996), Birgé (1997), Berlinet et al. (1998) and Meyer (2001) among others). Similarly, conditional mode estimation is typically carried out by conditional density estimation via different nonparametric methods (Gasser et al. (1998), Hall and Huang (2001) and Hall et al. (2001), Brunner (1992), Ho (2006), Dunson et al. (2007) among others). However, these nonparametric conditional density based mode regression models are not a direct estimation of the conditional mode. The problem with these methods is twofold: the estimation of the conditional density may suffer from the well-known “curse of dimensionality” and, it is hard to describe and interpret the estimated conditional mode in terms of predictors or covariates. Direct inference for mode regression was explored by Lee first in 1989 and then in 1993 (Lee,1989,1993). However, it has not been well-applied due to lack of proper inference tools. Recently, Kemp and Silva(2012) relaxed Lee’s restriction on truncated dependent variables and employed alternative kernel estimation. However, their regression coefficient estimation has slow convergence rate, involves bandwidth selection and provides only approximate normal confidence intervals. Moreover, direct Bayesian method for mode regression is not available but there is clear practical motivation from this perspective.

In this paper we introduce a fully Bayesian framework for direct mode regression inference by using three approaches: a parametric Bayesian method, a nonparametric Bayesian method and an empirical likelihood based Bayesian method. The remainder of the paper is organized as

follows. Section 2 introduces the three approaches, describes the theoretical and computational framework of these methods and gives their mathematical justification. In Section 3 we illustrate the proposed methods through two simulated case-studies and a real example. We conclude with a short discussion in Section 4.

2 Bayesian mode regression

Consider an arbitrary random variable Z , let $F_Z(z)$ be the distribution of Z with density $f_Z(z)$ and let $K(Z; \cdot)$ be the *step-loss function* (Manski, 1991) such as,

$$K(Z; \mu) = I\left[\frac{|Z - \mu|}{\sigma} \geq 1\right], \quad (1)$$

with $\sigma > 0$ and $I[A]$ being the indicator function of event A . If $f_Z(z)$ is symmetric about μ or if μ is the middle value of the interval of length 2σ that captures the most probability under $F_Z(z)$ then $\hat{\mu} = \operatorname{argmin}_{\mu} E\{L(Z; \mu)\}$ is the mode of Z . Lee (1989) introduced mode regression, or the conditional mode of y given \mathbf{x} , as $\operatorname{mode}(y|\mathbf{x}) = \mathbf{x}'\boldsymbol{\beta}$ based on the loss function $K(y; \mathbf{x}'\boldsymbol{\beta})$, where $\boldsymbol{\beta}$ is the regression parameter. That is, given a sample $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ from (\mathbf{x}, y) , when σ approaches 0, the parameter $\boldsymbol{\beta}$ in the conditional model of $y|\mathbf{x}$ is estimated by

$$\hat{\boldsymbol{\beta}} = \operatorname{argmin}_{\boldsymbol{\beta}} \sum_{i=1}^n K(y_i; \mathbf{x}_i'\boldsymbol{\beta}) \quad (2)$$

2.1 Parametric Bayesian method

The conditional mode denoted as $\operatorname{mode}(y|\mathbf{x}) = \mathbf{x}'\boldsymbol{\beta}$ can be reformulated as a standard regression model

$$y = \mathbf{x}'\boldsymbol{\beta} + \epsilon \quad (3)$$

with zero mode for the error term ϵ .

Given a sample $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ from (\mathbf{x}, y) , note that $\hat{\boldsymbol{\beta}} = \operatorname{argmax}_{\boldsymbol{\beta}} \sum_{i=1}^n I[|y_i - \mathbf{x}_i'\boldsymbol{\beta}| \leq \sigma]$. That is, $\hat{\boldsymbol{\beta}}$ in (2) can be regarded as the maximum likelihood estimates of the “working” likelihood function

$$L(y|\boldsymbol{\beta}) \propto \sigma^{-n} \sum_{i=1}^n I(|y_i - \mathbf{x}_i'\boldsymbol{\beta}| \leq \sigma). \quad (4)$$

Therefore, the Bayesian mode regression estimates, denoted as $\hat{\boldsymbol{\beta}}_B$, can be obtained using the posterior distribution of $\boldsymbol{\beta}$,

$$\pi(\boldsymbol{\beta}|y) \propto L(y|\boldsymbol{\beta}) \pi(\boldsymbol{\beta}), \quad (5)$$

where $\pi(\beta)$ is the prior distribution of β . Although a standard conjugate prior distribution is not available for the mode regression formulation, MCMC methods may be used for extracting the posterior distributions of β .

2.2 Consistency and asymptotic normality

The classical mode regression parameter estimator $\hat{\beta}$ of Lee (1989,1993)

$$\hat{\beta} = \operatorname{argmax}_{\beta} \frac{1}{n} \sum_{i=1}^n I[|y_i - \mathbf{x}'_i \beta| \leq \sigma].$$

is known to be consistent. According to White (1982), any posterior estimator $\hat{\beta}_B$ from the likelihood function (4) with a flat prior, even if misspecified, is still consistent, in the sense of minimization of the Kullback-Leibler distance between the true distribution and the parametric family to which the approximation belongs or in the spirit of the quasi-maximum likelihood estimator of White (1982).

Further, although Bayesian inference does not require large sample theory we provide evidence that the posterior distribution obtained via the proposed Bayesian approach, under certain regularity conditions, is asymptotically normal when the sample size increases. This is the same as the classical mode regression estimator $\hat{\beta}$ whose asymptotic normality was derived under a special case of “M-estimators” (Huber (1973), Lee (1993)).

In fact, let $I(\beta) = E[\{\frac{\partial}{\partial \beta} \log f(y|\beta)\}^2]$ be the total Fisher information in the data and define $I_n(\hat{\beta}_B) = I(\beta)|_{\beta=\hat{\beta}_B}$. Under certain regularity conditions, Taylor power series expansions of the logarithm of the posterior distribution leads to

$$\log \pi(\beta|y) = \log \pi(\hat{\beta}_B|y) - \frac{1}{2}(\beta - \hat{\beta}_B)^T I_n(\hat{\beta}_B)(\beta - \hat{\beta}_B),$$

hence,

$$\begin{aligned} \pi(\beta|y) &\approx \pi(\hat{\beta}_B|y) \exp[-\frac{1}{2}(\beta - \hat{\beta}_B)^T I_n(\hat{\beta}_B)(\beta - \hat{\beta}_B)] \\ &\propto \exp[-\frac{1}{2}(\beta - \hat{\beta}_B)^T I_n(\hat{\beta}_B)(\beta - \hat{\beta}_B)], \end{aligned}$$

which is the kernel of a $N_p(\beta|\hat{\beta}_B, I_n^{-1}(\hat{\beta}_B))$ density. This implies that $\sqrt{n}(\hat{\beta}_B - \beta) \sim N(0, I_n^{-1}(\hat{\beta}_B))$, where $I_n^{-1}(\hat{\beta}_B)$ is specified below, which has the same form as the asymptotic distribution of “M-estimators”. On the other hand, the likelihood function associated with Bayesian mode regression can also be formulated as

$$L(\beta) = \frac{e^n}{(2\sigma)^n} \prod_{i=1}^n \exp(-I[|y_i - \mathbf{x}'_i \beta| \leq \sigma]),$$

which is a likelihood function based on the uniform probability density

$$f_\sigma(u) = \frac{e}{2\sigma} \exp(-I[|u - \mu| \leq \sigma]), \quad (6)$$

for a window parameter $\sigma > 0$. Then, under a flat prior, from

$$f(y|\mathbf{x}, \boldsymbol{\beta}, \sigma) = \frac{e}{2\sigma} \exp[-I[|y - \mathbf{x}'\boldsymbol{\beta}| \leq \sigma]],$$

we have

$$\begin{aligned} \log f(y|\mathbf{x}, \boldsymbol{\beta}, \sigma) &= 1 - \log(2\sigma) - I[|y - \mathbf{x}'\boldsymbol{\beta}| \leq \sigma], \\ \frac{\partial}{\partial \boldsymbol{\beta}} \log f(y|\mathbf{x}, \boldsymbol{\beta}, \sigma) &= -I[|y - \mathbf{x}'\boldsymbol{\beta}| \leq \sigma], \\ I_n(\hat{\boldsymbol{\beta}}_B) &= E\left[\left\{\frac{\partial}{\partial \boldsymbol{\beta}} \log f(y|\boldsymbol{\beta})\right\}^2\right]_{\hat{\boldsymbol{\beta}}_B} = E\{I[|y - \mathbf{x}'\boldsymbol{\beta}| \leq \sigma]\}_{\hat{\boldsymbol{\beta}}_B}. \end{aligned}$$

Thus, the asymptotic justification of using the proposed “working” likelihood for parametric Bayesian mode regression is fully outlined.

2.3 The estimation of covariance matrix of classical estimates

Under the classical approaches of Lee (1989, 1993) and Kemp and Silva (2012), the covariance matrix, $cov\{\hat{\boldsymbol{\beta}}\}$ of the classical estimator $\hat{\boldsymbol{\beta}}$ and its inverse are often required but difficult to estimate or compute numerically, especially under small or moderate samples. A by-product of the proposed Bayesian approach is that using the MCMC posterior sample leads to a natural and efficient estimation of $cov\{\hat{\boldsymbol{\beta}}\}$ and other asymptotic quantities of $\hat{\boldsymbol{\beta}}$.

In fact, a MCMC scheme constructs a Markov chain whose equilibrium distribution is the joint posterior, $p(\boldsymbol{\beta}|data)$. After running the Markov chain for a burn-in period, one obtains samples from the limiting distribution, provided that the Markov chain has converged. Given that the chain has converged, the frequency of appearance of the parameters in the Markov chain represents their posterior distribution. An informative full density distribution of the model parameters is readily obtained rather than a single point estimate as in the classical approach.

When a Markov chain, S , is drawn from the posterior distribution, $p(\boldsymbol{\beta}|data)$: $S = (\boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, \dots, \boldsymbol{\beta}^{(N)})$, where N is the number of draws after burn-in, a consistent estimate of the inverse of the covariance matrix $cov\{\hat{\boldsymbol{\beta}}\}$ can be obtained by multiplying by N the variance-covariance matrix of this MCMC sequence (Chernozhukov and Hong, 2003).

2.4 Prior selection and proper posteriors

In this section we demonstrate that in the absence of any realistic information one could use improper uniform prior distributions for all the components of $\boldsymbol{\beta}$ as such a choice yields a proper joint posterior and then we address the important issue of specifying a prior for parameter σ .

Below, we show that if we choose the prior of β to be improper uniform, then the resulting joint posterior distribution will be proper.

Theorem 1: If the likelihood function is given by (4) and $\pi(\beta) \propto 1$, then the posterior distribution of β , $\pi(\beta|\mathbf{y})$, will be a proper distribution. In other words

$$0 < \int \pi(\beta|\mathbf{y}) d\beta < \infty,$$

or, equivalently,

$$0 < \int L(\mathbf{y}|\beta) \pi(\beta) d\beta < \infty.$$

The proof can be found in Appendix A.

In practice one usually assumes that the components of β have independent improper uniform prior distributions which is a special case of the above theorem.

Next we address the issue of determining a suitable prior for the parameter σ . The aim is to manage to determine a value that is neither too small nor too large to prevent underutilization of data with $|y_i - x'_i b| < \sigma$ or with $|y_i - x'_i b| > \sigma$ respectively. Lee (1989, 1993) suggested some possible methods for determining the value of σ (including trial and error and bootstrapping methods).

In this work we apply a $\text{Uniform}(w_1, w_2)$ prior on σ , where w_i can be determined using one or more of the following three rules-of-thumb, depending on the assumption for the underlying distribution.

- The empirical rule, which states that, given a symmetric distribution, approximately 99.7% of the data values fall within three standard deviations (sd) of the mean, therefore, $w_i = 3 * \hat{sd}$;
- Chebyshev's Theorem, which is true for any sample set no matter what the distribution is, and states that at least 93.75% of the data values fall within four standard deviations of the mean, therefore, $w_i = 4 * \hat{sd}$;
- Variations of Silverman's plug-in estimate for the bandwidth (Silverman, 1986), a simple formula for w_i that depends on the sample size n and the sample standard deviation sd , given by $w_i = 1.3643\delta n^{-0.2}[\min(\hat{sd}, IQR/1.349)]$ where IQR is the sample inter quantile range and $\delta = 1.3510$ for a uniform kernel. This formula assumes data which is normally distributed and uses $IQR/1.349$ as an alternative estimate of w_i that protects against outliers. These plug-in estimates for w_i work well in practice, especially for symmetric unimodal densities even if the data is not normally distributed. Alternatively, $IQR/1.349$ can be replaced by $1.4826 * MAD$ to cover data with large number of outliers.

However, the choice of a suitable prior for σ can be difficult in practice, therefore with the aim of developing a more flexible model, in the following section, we relax the distributional assumption on the prior for σ using a Dirichlet process prior. This leads to a flexible nonparametric mixture model. The method is nonparametric in the sense that we do not assume that the prior belongs to any fixed class of distributions.

2.5 Nonparametric Bayesian method

In this section, we formulate a nonparametric Bayesian mode regression model to avoid critical dependence on the mode uniform distribution assumption thus to address the issue of misspecification that may arise under the parametric Bayesian method.

A density $f(\cdot)$ on \mathbb{R}^+ is non-increasing if and only if there exists a distribution function G such that $f(x|G) = \int \sigma^{-1} I_{[0 < x < \sigma]} dG(\sigma)$ (Feller, 1971). Therefore, any unknown density $f(\cdot)$ (with mode θ), symmetric or not, can be represented as a scale mixture of symmetric uniform distributions, that is

$$f(x|\theta, G) = \int \frac{1}{2\sigma} I_{[-\sigma < x - \theta < \sigma]} dG(\sigma), \quad (7)$$

where G is the mixing distribution supported on \mathbb{R}^+ .

Then a nonparametric Bayesian mode regression model can be expressed in the hierarchical form

$$\begin{aligned} y_i | \beta, \sigma_i &\stackrel{ind}{\sim} f(y_i - x'_i \beta; \sigma_i), i = 1 \cdots n \\ \sigma_i | G &\stackrel{iid}{\sim} G, i = 1 \cdots n \\ G | M, d &\sim DP(M, G_0(\cdot, d)) \\ \beta, M, d &\sim p(\beta), p(M), p(d), \end{aligned} \quad (8)$$

where, G is the mixing distribution, with base distribution G_0 and concentration parameter M and $f(y_i - x'_i \beta; \sigma_i) = \frac{1}{2\sigma} I_{[-\sigma < y_i - x'_i \beta < \sigma]}$ is the density of a uniform distribution on $(-\sigma, \sigma)$.

We take a uniform distribution as the base distribution, G_0 , uniform prior for M and we choose non-informative Normal priors for all the components of β .

2.6 Empirical likelihood based Bayesian method

In addition to parametric and nonparametric likelihood, an empirical likelihood based method could be an alternative for Bayesian mode regression. To derive an empirical likelihood for mode regression we begin with notations and a moment restriction. Lee (1993) generalized the mode regression estimator of Lee (1989), $\hat{\beta} = \operatorname{argmin}_{\beta} E\{L(Y - \mathbf{x}'\beta)\}$, by using the rectangular kernel

$L(Y; \mu) = \{(\sigma^2 - (Y - \mu)^2)I[|Y - \mu| < \sigma]\}$. Therefore, the moment restriction for the empirical likelihood can be obtained by the derivative $\frac{\partial}{\partial \mu} L(Y; \mu) = 2(Y - \mu)I[|Y - \mu| < \sigma]$. Let $l(\mu)$ be the ‘derivative’ of $L(.; \mu)$, then the mode, μ , of Y satisfies the moment restriction $E(l(\mu)) = 0$, where $l(\mu) = (Y - \mu)I[|Y - \mu| < \sigma]$.

Under an empirical likelihood for mode regression $\mu = \mathbf{x}'\boldsymbol{\beta}$, thus for any proposed $\boldsymbol{\beta}$ to estimate the true p dimensional $\boldsymbol{\beta}_0$ via empirical likelihood, we use the vector estimating functions $g(X, Y, \boldsymbol{\beta})$ with component $g_j(X, Y, \boldsymbol{\beta}) = l(X, Y, \boldsymbol{\beta}) X_j$ for $j = 1, \dots, p$. Then, the profile empirical likelihood ratio is given by

$$\mathfrak{R}(\boldsymbol{\beta}) = \max \left\{ \prod_{i=1}^n (n p_i) \mid \sum_{i=1}^n p_i g(X_i, Y_i, \boldsymbol{\beta}) = 0, p_i \geq 0, \sum_{i=1}^n p_i = 1 \right\}.$$

By a standard Lagrange multiplier argument we have

$$\mathfrak{R}(\boldsymbol{\beta}) = \prod_{i=1}^n \{n p_i(\boldsymbol{\beta})\},$$

with the weights $p_i(\boldsymbol{\beta}) = \frac{1}{n(1 + \lambda' g(X_i, Y_i, \boldsymbol{\beta}))}$, where the Lagrange multiplier λ satisfies

$$\sum_{i=1}^n \frac{g(X_i, Y_i, \boldsymbol{\beta})}{1 + \lambda' g(X_i, Y_i, \boldsymbol{\beta})} = 0.$$

According to Qin and Lawless (1994), among others, the existence and uniqueness of λ are guaranteed when the following two conditions are satisfied: (1) zero belongs the convex hull of $\{g(X_i, Y_i, \boldsymbol{\beta}), i = 1, \dots, n\}$ and (2) the matrix $\sum_{i=1}^n \{g(X_i, Y_i, \boldsymbol{\beta})g(X_i, Y_i, \boldsymbol{\beta})'\}$ is positive definite.

Under Bayesian inference we consider the empirical likelihood function $\mathfrak{R}(\boldsymbol{\beta})/n^n = \prod_{i=1}^n \{p_i(\boldsymbol{\beta})\}$, which can be combined with a prior specification $\pi(\boldsymbol{\beta})$ on the parameter $\boldsymbol{\beta}$ to obtain the posterior distribution

$$\pi(\boldsymbol{\beta} | \text{data}) \propto \pi(\boldsymbol{\beta}) \mathfrak{R}(\boldsymbol{\beta}).$$

2.7 Asymptotic properties of Bayesian empirical likelihood

Before studying the asymptotic normality of the empirical likelihood based Bayesian mode regression parameter estimates, we should provide the consistency of the empirical likelihood estimator, which is a necessary condition for the asymptotic normality of the posterior. As the criterion function $g(X, Y, \boldsymbol{\beta})$ results in non-smooth estimating equations we employ a similar method to the one use by Molanes Lopez et al. (2009), among others, to derive our asymptotic results.

Let $\hat{\boldsymbol{\beta}} = \operatorname{argmax}_{\boldsymbol{\beta}} \mathfrak{R}(\boldsymbol{\beta})$ be the empirical likelihood estimates in a compact set of parameter space which contains the true parameter $\boldsymbol{\beta}_0$. Then note that our criterion function $g(X, Y, \boldsymbol{\beta})$ can be regarded as a special case of M-estimators as discussed in Chapter 5 of Van der Vaart (1998) and

satisfies the conditions of theorem 5.7 in the book. Under some regular conditions such as uniformly continuous and bounded imposed on the marginal distribution of X and conditional distribution of Y given X , and assume the matrix $E\{g(X, Y, \beta) g(X, Y, \beta)'\} > 0$, then $E\{g(X, Y, \beta)\}$ is sufficiently smooth in a compact set of parameter space which contains β_0 , so the consistency condition C_3 of Molanes Lopez et al. (2009) holds, that is, the consistency of β 's empirical likelihood estimates is established.

The asymptotic normality of the posterior distribution $\pi(\beta|data)$ could be established using the fact that the empirical log-likelihood ratio for β is well approximated by certain quadratics in the sense of Lemma 6 of Molanes Lopez et al. (2009) so that,

$$\Gamma_n(\beta) \equiv -n^{-1} \sum_{i=1}^n \log(1 + \lambda^T g(X_i, Y_i, \beta))$$

$$= -\frac{1}{2}(\beta - \beta_0)' V_{12} V_{11}^{-1} V_{12} (\beta - \beta_0) + n^{-1/2} (\beta - \beta_0)' V_{12}' V_{11}^{-1} W_n - \frac{1}{2} n^{-1} W_n' V_{11}^{-1} W_n + o_P(n^{-1}),$$

with matrices $V_{11} = (E\{g_j(X, Y, \beta) g_k(X, Y, \beta)'\})$, $V_{12} = -\frac{\partial}{\partial \beta_k} E\{g_j(X, Y, \beta)\}$, and vector $W_n = n^{-1/2} \sum_{i=1}^n g(X_i, Y_i, \beta)$.

Then from $\log \Re(\beta) = n\Gamma_n(\beta)$ we have the posterior

$$\pi(\beta|data) = \pi(\beta) \Re(\beta) \propto \exp\left\{-\frac{1}{2}(\beta - \hat{\beta})' I_n (\beta - \hat{\beta}) + O_p(1)\right\},$$

where $I_n = nV_{12}' V_{11}^{-1} V_{12}$ and $\hat{\beta}$ is the empirical likelihood estimate.

3 Numerical experiments

In this section we demonstrate our approach to Bayesian mode regression through two simulated and one real examples. The real example is based on the Western Electric Workers (WECO) dataset and investigates how the worker's gender, pre-employment test result and education, can affect productivity.

3.1 Simulation example 1

We consider a simulated data from the model

$$y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \tag{9}$$

where $x_i \sim N(0, 1)$ and $i = 1, \dots, n$, with $n = 50, 100, 200$. We set $\beta = (1, 2)$ and consider the following three specifications for the model error ϵ :

- Case 1: the standard normal distribution, $\epsilon_i \sim N(0, 1)$ - a symmetric error distribution;

- Case 2: a Fisher’s Z distribution, $\epsilon_i \sim 1/2\log Z$ with $Z \sim F_{2,2}$ - a skewed error distribution;
- Case 3: a normal distribution with normally distributed outliers (contaminants) centered at twice the distance between the true mode and the 99th percentile of the original normal distribution and accounting for 20% of the total data points, $\epsilon_i \sim 0.80N(0, \frac{1}{4}) + 0.20N(2.5, \frac{1}{4})$ (Hedges and Shah, 2003) - an asymmetric error distribution.

We fit parametric Bayesian mode regression (labeled PBMR) for all the cases above. Then for demonstration and comparison purposes we fit empirical likelihood based Bayesian mode regression (labeled ELBMR) for case 2 and nonparametric Bayesian mode regression (labeled NBMR) for case 3.

For PBMR and ELBMR, we chose independent improper uniform priors for all the components of β and we simulated realizations from the posterior distributions by means of a single-component Metropolis-Hastings algorithm. Each of the parameters was updated using a random-walk Metropolis algorithm with a Gaussian proposal density centered at the current state of the chain. The variance of the proposal density was determined to provide an acceptance rate close to the optimal acceptance rate as defined in Roberts and Rosenthal (2001). Convergence was assessed using time series plots and the R package *boa* (Smith, 2007). The estimates are posterior means using 10,000 iterations of the MCMC sampler (after 10,000 burn-in iterations).

The estimates for NBMR were obtained by fitting a truncated Dirichlet Process (DP) mixture model, which leads to a computationally straightforward approximation and can be easily implemented in the freely available WinBUGS software. Two parallel chains of equal length with different initial values were run for the model. The results were based on 10,000 iterations which followed a burn-in period of 40,000 for each chain.

Table 1 compares the posterior means with the true values of β_0 and β_1 and also gives standard deviations and 95% credible intervals for each of the models considered in this example .

As expected, the PBRM works well as all the absolute biases for the estimated parameters turn out to be in the range $[0.01, 0.26]$. Furthermore, under both ELBMR and NBRM, the true values for both β_0 and β_1 are recovered successfully indicating that the methods also work well. However, it should be noted that the standard deviations for both parameters are smaller than in the PBMR, giving shorter confidence intervals.

Figure 1 exhibits the empirical samples from the joint posterior distributions of the PBMR parameters, which were obtained using the output of the MCMC sampler for the regression parameters $\hat{\beta}_0$ and $\hat{\beta}_1$. These samples can be used to obtain a consistent estimator of the covariance or correlation structure of the parameter estimators, which is difficult to estimate under the classical


Table 1: Simulation Example 1: True parameter values (T.V.) and their posterior means, standard deviations (S.D.) and 95% credible intervals

| | | PBMR | | | | | | ELBMR | | NBMR | |
|-----|---------|------------|-----------|------------|-----------|------------|-----------|-------------|-------------|-------------|-------------|
| | | Normal | | Skewed | | Asymmetric | | Skewed | | Asymmetric | |
| n | | β_0 | β_1 | β_0 | β_1 | β_0 | β_1 | β_0 | β_1 | β_0 | β_1 |
| 50 | T.V | 1 | 2 | 1 | 2 | 1 | 2 | 1 | 2 | 1 | 2 |
| | Mean | 0.92 | 2.00 | 1.07 | 2.01 | 0.96 | 2.02 | 1.01 | 2.00 | 1.09 | 1.94 |
| | S.D. | 0.78 | 0.77 | 0.78 | 0.49 | 0.34 | 0.24 | 0.01 | 0.01 | 0.24 | 0.19 |
| | 95% HPD | (-0.6,2.1) | (0.5,3.3) | (-0.3,2.6) | (1.2,3.1) | (0.4,1.7) | (1.6,2.5) | (0.99,1.02) | (1.99,2.01) | (0.7,1.5) | (1.5,2.3) |
| 100 | T.V | 1 | 2 | 1 | 2 | 1 | 2 | 1 | 2 | 1 | 2 |
| | Mean | 1.01 | 2.10 | 0.95 | 1.89 | 1.06 | 1.94 | 1.01 | 2.00 | 1.06 | 2.00 |
| | S.D. | 0.18 | 0.25 | 0.52 | 0.37 | 0.98 | 0.76 | 0.01 | 2 | 0.14 | 0.12 |
| | 95% HPD | (0.6,1.3) | (1.6,2.6) | (0.0,1.9) | (1.2,2.6) | (-0.7,2.9) | (0.5,3.3) | (0.99,1.02) | (1.99,2.01) | (0.8,1.3) | (1.8,2.2) |
| 200 | T.V | 1 | 2 | 1 | 2 | 1 | 2 | 1 | 2 | 1 | 2 |
| | Mean | 1.26 | 1.99 | 1.00 | 1.99 | 1.06 | 1.96 | 1.01 | 2.00 | 1.04 | 1.91 |
| | S.D. | 0.86 | 0.52 | 1.29 | 0.75 | 0.82 | 0.42 | 0.01 | 0.01 | 0.07 | 0.06 |
| | 95% HPD | (-0.5,2.8) | (0.9,3.0) | (-1.3,3.5) | (0.6,3.3) | (-0.4,2.6) | (1.2,2.7) | (0.99,1.02) | (1.99,2.01) | (0.92,1.19) | (1.78,2.03) |


approach. For example in case (a), with sample size n=100,

$$\widehat{Cov}\begin{pmatrix} \hat{\beta}_0 \\ \hat{\beta}_1 \end{pmatrix} = \begin{pmatrix} 3 & -1 \\ -1 & 6 \end{pmatrix}$$


Figure 1: Plots showing the empirical samples from the joint distributions of mode regression parameters



(a) Symmetric error



(b) Skewed error



(c) Asymmetric error

3.2 Simulation example 2

In this section we present the results of a second simulation example with the aim of comparing the performance of our approach with the classical mode regression approach. Specifically, we replicate the simulation study in Kemp and Silva(2012), but only for a sample of size 250, and

compare their results with the results obtained under our Bayesian mode regression approach.

Simulation data are generated by the simple linear model

$$y_i = \beta_0 + \beta_1 x_i + (1 + vx_i)\epsilon_i, \quad (10)$$

where x_i are generated from a $\chi^2_{(3)}$ distribution, scaled to have variance 1, and ϵ_i are generated as independent draws from a re-scaled log-gamma random variable,

$$\epsilon_i = -\lambda \ln(Z_i) \quad (11)$$

where Z follows a gamma distribution with mean 1 and scale parameter $\frac{1}{\alpha}$, to ensure that ϵ_i has zero mode. Furthermore, we set $\lambda = [(1 + 2E(x_i)v + E(x_i^2)v^2)\psi(\alpha)]^{-1}$ to ensure that the unconditional variance of the error $(1 + vx_i)$ is equal to one.

The study was performed for $\alpha \in \{0.05, 5\}$ and for $v \in \{0, 2\}$. Table 2 compares the 95% credible intervals for the estimates obtained under PBMR and NBMR with the 95% confidence intervals for the estimates under the two classical mode regression models: Mode 1.6 and Mode 0.8. Mode 1.6 and Mode 0.8 correspond to $k = 1.6$ and $k = 0.8$ respectively in the bandwidth selection rule, $\text{bandwidth} = k \text{mad} n^{-0.143}$, with mad = the median of the absolute deviation from the median of ordinary least squares regression residuals.

Table 2: Simulation Example 2: Comparison between Classical and Bayesian approach for mode regression

| | | | PBMR | NBMR | Mode 1.6 | Mode 0.8 |
|----------|---|-----------|--------------|--------------|---------------|---------------|
| α | n | | 95% HPD | 95% HPD | 95% CI | 95% CI |
| 5.00 | 0 | β_0 | (-0.37,0.29) | (-0.21,0.36) | (-0.31, 0.41) | (-0.69, 0.75) |
| | | β_1 | (0.82,1.28) | (0.89,1.32) | (0.77, 1.24) | (0.56,1.45) |
| | 2 | β_0 | (-0.06,0.07) | (-0.03,0.21) | (-0.15,0.23) | (-0.25,0.29) |
| | | β_1 | (0.99,1.14) | (0.80,1.22) | (0.63,1.37) | (0.48,1.53) |
| 0.05 | 0 | β_0 | (0.00, 0.14) | (-0.03,0.07) | (0.12,0.42) | (-0.09,0.35) |
| | | β_1 | (0.95,1.13) | (0.95,1.06) | (0.90,1.11) | (0.87,1.17) |
| | 2 | β_0 | (0.02,0.08) | (0.04,0.09) | (0.09,0.29) | (0.01,0.21) |
| | | β_1 | (0.99,1.08) | (0.97,1.04) | (0.91,1.19) | (0.85,1.19) |

The results of the analysis suggest that the Bayesian mode regression estimates are strong competitors of the classical mode regression estimates since in almost all the examples both PBMR

¹ $\psi(\cdot)$ is the trigamma function

and NBMR estimators outperform the two classical estimators.

Finally, as also evident from Kemp and Silva(2012), the selection of the value/prior for σ plays an important role on the precision of the parameters, an issue that is less restrictive under NBMR.

3.3 Productivity of Western Electric Workers - WECO

To illustrate the applicability of our approach we consider a model for predicting the productivity of newly hired Electric workers in a manufacturing firm. Productivity (y_i) was modeled as a function of a gender indicator (sex_i), the score on a physical dexterity exam administered prior to employment (dex_i) and the years of education (lex_i).

$$y_i = \beta_0 + \beta_1 sex_i + \beta_2 dex_i + \beta_3 lex_i + \beta_4 lex_i^2 + \epsilon_i \quad (12)$$

The data come originally from the study of Klein et al. (1991), but have been modified over the years to heighten the pedagogical impact. Figure 2 presents the density plot for productivity which is unimodal and almost symmetric (skewness =0.069).

While the productivity levels range from 10.5 to 19.1, one is interested in how the typical productivity level is affected by the model covariates. To estimate this effect we apply our PBMR model to estimate the model parameters, $\beta_0, \beta_1, \beta_2, \beta_3$ and β_4 . The output was obtained by running the sampler for 50,000 cycles after a burn-in of 100,000, to ensure convergence and mixing. Table 3 summarizes the results.

The results indicate that on average the mode productivity level of a female worker, who scores zero on her physical dexterity exam and has zero years of education is 4.93 units. Furthermore, it can be concluded that on average the most frequent productivity level is lower for a male worker, while it is higher for workers with a higher exam score. Finally, it is deduced that an additional year of education contributes positively to the level of mode productivity.

Given that under the PBMR a relatively wide credible interval is obtained for the some model parameters we also fit a NBMR to the WECO dataset. Again, two parallel chains of equal length with different initial values were run for the model. The results were based on 20,000 iterations which followed a burn-in period of 50,000 iterations for each chain. As illustrated in Table 3, the results obtained under the NBMR are similar to the results obtained under PBRM, but now the confidence intervals are much smaller.

Figure 2: Density plot for WECO data




Table 3: Model parameters and their estimated posterior means, standard deviations (S.D.) and 95% Credible intervals for the WECO data

| | PBMR | | | NBMR | | |
|-----------|-------|------|---------------|-------|-------|----------------|
| Parameter | Mean | S.D. | 95% HPD | Mean | S.D. | 95% HPD |
| β_0 | 4.93 | 8.13 | (-12.6,19.8) | 4.10 | 1.11 | (2.56, 6.52) |
| β_1 | -0.71 | 0.46 | (-1.58,0.21) | -0.84 | 0.08 | (-1.03,-0.73) |
| β_2 | 0.12 | 0.03 | (0.06,0.18) | 0.12 | 0.005 | (0.11,0.12) |
| β_3 | 0.87 | 1.27 | (-1.44,3.51) | 1.08 | 0.18 | (0.69,1.37) |
| β_4 | -0.04 | 0.05 | (-0.14,0.06) | -0.05 | 0.008 | (-0.06,-0.03) |

4 Conclusions

In this paper we introduce a novel Bayesian mode regression framework which includes three approaches: a parametric method, a nonparametric method and an empirical likelihood based method, as in the area of mode regression, there is no literature from a Bayesian perspective. We demonstrate that our estimates are consistent and asymptotically normal under rather standard conditions, even under misspecification of the likelihood function. The approaches are easy to

implement and have proper inference tools as well as credible intervals available irrespective of the sample size. The numerical studies suggest that the proposed Bayesian mode regression estimates are strong competitors of the classical mode regression estimates.

Appendix A

Proof of theorem 1

For any $\sigma > 0$ and $m > p$, the moments of posterior distribution is given by

$$E[|\boldsymbol{\beta}|^\gamma | \sigma, \mathbf{y}] = \int \prod_{j=0}^p |\beta_j|^{r_j} \frac{e^n}{(2\sigma)^n} \sum_{i=1}^n \exp[-I[|y_i - x'_i \boldsymbol{\beta}| < \sigma]] d\boldsymbol{\beta}.$$

Noting that $\sum_{i=1}^n \exp[-I[|y_i - x'_i \boldsymbol{\beta}| < \sigma]]$ is always a constant whether $|y_i - x'_i \boldsymbol{\beta}| < \sigma$ or not ($i = 1, \dots, n$). Suppose that the coefficient matrix $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p)$ of mode regression equations $y_i = \mathbf{x}_i^T \boldsymbol{\beta} + \epsilon_i$ is a full rank matrix with rank p , then there is a subset of p constrains $|y_i - x'_i \boldsymbol{\beta}| < \sigma$ ($i = 1, \dots, n$) to provide $0 < |\beta_j| < B_j < \infty$ ($j = 0, 1, \dots, p-1$), even if some of $|y_i - x'_i \boldsymbol{\beta}| < \sigma$ are true and some are not. Therefore,

$$E[|\boldsymbol{\beta}|^\gamma | \sigma, \mathbf{y}] = \text{const.} \int_{-B_0}^{B_0} \int_{-B_1}^{B_1} \dots \int_{-B_p}^{B_p} \prod_{j=0}^p |\beta_j|^{r_j} d\boldsymbol{\beta},$$

which is finite.

References

- Berlinet, A., I. Vajda, and E. Van der Meulen (1998). About the asymptotic accuracy of barron density estimates. *Information Theory, IEEE Transactions* 44(3), 999–1009.
- Bickel, P. and J. Fan (1996). Some problems on the estimation of unimodal densities. *Statistica Sinica* 6, 23–46.
- Birgé, L. (1997). Estimation of unimodal densities without smoothness assumptions. *The Annals of Statistics* 25(3), 970–981.
- Brunner, L. (1992). Bayesian nonparametric methods for data from a unimodal density. *Statistics & Probability letters* 14(3), 195–199.
- Chernozhukov, V. and H. Hong (2003). An mcmc approach to classical estimation. *Journal of Econometrics* 115(2), 293–346.

- Dunson, D., N. Pillai, and J. Park (2007). Bayesian density regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(2), 163–183.
- Eddy, W. (1980). Optimum kernel estimators of the mode. *The Annals of Statistics* 8(4), 870–882.
- Feller, W. (1971). *An introduction to probability theory and its applications*, Volume 2. Wiley-New York.
- Gasser, T., P. Hall, and B. Presnell (1998). Nonparametric estimation of the mode of a distribution of random curves. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 60(4), 681–691.
- Grenander, U. (1965). Some direct estimates of the mode. *The Annals of Mathematical Statistics* 36(1), 131–138.
- Hall, P. and L. Huang (2001). Nonparametric kernel regression subject to monotonicity constraints. *The Annals of Statistics* 29(3), 624–647.
- Hall, P., L. Peng, and C. Rau (2001). Local likelihood tracking of fault lines and boundaries. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(3), 569–582.
- Heckman, D., D. Geiser, B. Eidell, R. Stauffer, N. Kardos, and S. Hedges (2001). Molecular evidence for the early colonization of land by fungi and plants. *Science* 293(5532), 1129.
- Hedges, S. and P. Shah (2003). Comparison of mode estimation methods and application in molecular clock analysis. *BMC Bioinformatics* 4(1), 31.
- Huber, P. (1973). Robust regression: asymptotics, conjectures and monte carlo. *The Annals of Statistics* 1(5), 799–821.
- Kemp, G. and J. Silva (2012). Regression towards the mode. *Journal of Econometrics*, under revision.
- Klein, R., R. Spady, and A. Weiss (1991). Factors affecting the output and quit propensities of production workers. *Review of Economic Studies* 58, 929–954.
- Kumar, S. and S. Hedges (1998). A molecular timescale for vertebrate evolution. *Nature* 392(6679), 917–920.
- Lee, M. (1989). Mode regression. *Journal of Econometrics* 42(3), 337–349.
- Lee, M. (1993). Quadratic mode regression. *Journal of Econometrics* 57(1-3), 1–19.

- Manski, C. (1991). Regression. *Journal of Economic Literature* 29(1), 34–50.
- Markov, H., T. Valtchev, J. Borissova, and V. Golev (1997). An algorithm to ”clean” close stellar companions. *Astronomy and Astrophysics Supplement Series* 122(1), 193–199.
- Meyer, M. (2001). An alternative unimodal density estimator with a consistent estimate of the mode. *Statistica Sinica* 11(4), 1159–1174.
- Molanes Lopez, E., I. Keilegom, and N. Veraverbeke (2009). Empirical likelihood for non-smooth criterion functions. *Scandinavian Journal of Statistics* 36(3), 413–432.
- Parzen, E. (1962). On estimation of a probability density function and mode. *The Annals of Mathematical Statistics* 33(3), 1065–1076.
- Qin, J. and J. Lawless (1994). Empirical likelihood and general estimating equations. *The Annals of Statistics*, 300–325.
- Roberts, G. and J. Rosenthal (2001). Optimal scaling for various metropolis-hastings algorithms. *Statistical Science* 16(4), 351–367.
- Silverman, B. (1986). *Density estimation for statistics and data analysis*, Volume 26. Chapman & Hall/CRC.
- Smith, B. (2007). boa: an r package for mcmc output convergence assessment and posterior inference. *Journal of Statistical Software* 21(11), 1–37.
- Van der Vaart, A. (1998). *Asymptotic statistics*. Cambridge Univ Pr.
- White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society* 50(1), 1–25.
- Yasukawa, K. (1926). On the probable error of the mode of skew frequency distributions. *Biometrika* 18(3/4), 263–292.

